

**EV355226571**

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

APPLICATION FOR LETTERS PATENT

**Kernel-Mode Audio Processing Modules**

Inventor(s):

Martin G. Puryear

ATTORNEY'S DOCKET NO. MS1-541USC1

## **RELATED APPLICATIONS**

This application is a continuation of U.S. Patent Application Serial No. 09/559,986, filed April 26, 2000, entitled "Kernel-Mode Audio Processing Modules" to Martin G. Puryear, which claims the benefit of U.S. Provisional Application No. 60/197,100, filed April 12, 2000, entitled "Extensible Kernel-Mode Audio Processing Architecture" to Martin G. Puryear.

## **TECHNICAL FIELD**

This invention relates to audio processing systems. More particularly, the invention relates to kernel-mode audio processing modules.

## **BACKGROUND OF THE INVENTION**

Musical performances have become a key component of electronic and multimedia products such as stand-alone video game devices, computer-based video games, computer-based slide show presentations, computer animation, and other similar products and applications. As a result, music generating devices and music playback devices are now tightly integrated into electronic and multimedia components.

Musical accompaniment for multimedia products can be provided in the form of digitized audio streams. While this format allows recording and accurate reproduction of non-synthesized sounds, it consumes a substantial amount of memory. As a result, the variety of music that can be provided using this approach is limited. Another disadvantage of this approach is that the stored music cannot be easily varied. For example, it is generally not possible to change a particular musical part, such as a bass part, without re-recording the entire musical stream.

1       Because of these disadvantages, it has become quite common to generate  
2 music based on a variety of data other than pre-recorded digital streams. For  
3 example, a particular musical piece might be represented as a sequence of discrete  
4 notes and other events corresponding generally to actions that might be performed  
5 by a keyboardist—such as pressing or releasing a key, pressing or releasing a  
6 sustain pedal, activating a pitch bend wheel, changing a volume level, changing a  
7 preset, etc. An event such as a note event is represented by some type of data  
8 structure that includes information about the note such as pitch, duration, volume,  
9 and timing. Music events such as these are typically stored in a sequence that  
10 roughly corresponds to the order in which the events occur. Rendering software  
11 retrieves each music event and examines it for relevant information such as timing  
12 information and information relating the particular device or “instrument” to  
13 which the music event applies. The rendering software then sends the music event  
14 to the appropriate device at the proper time, where it is rendered. The MIDI  
15 (Musical Instrument Digital Interface) standard is an example of a music  
16 generation standard or technique of this type, which represents a musical  
17 performance as a series of events.

18       Computing devices, such as many modern computer systems, allow MIDI  
19 data to be manipulated and/or rendered. These computing devices are frequently  
20 built based on an architecture employing multiple privilege levels, often referred  
21 to as user-mode and kernel-mode. Manipulation of the MIDI data is typically  
22 performed by one or more applications executing in user-mode, while the input of  
23 data from and output of data to hardware is typically managed by an operating  
24 system or a driver executing in kernel-mode.

1       Such a setup requires the MIDI data to be received by the driver or  
2 operating system executing in kernel-mode, transferred to the application  
3 executing in user-mode, manipulated by the application as needed in user-mode,  
4 and then transferred back to the operating system or driver executing in kernel-  
5 mode for rendering. Data transfers between kernel-mode and user-mode,  
6 however, can take a considerable and unpredictable amount of time. Lengthy  
7 delays can result in unacceptable latency, particularly for real-time audio  
8 playback, while unpredictability can result in an unacceptable amount of jitter in  
9 the audio data, resulting in unacceptable rendering of the audio data.

10       The invention described below addresses these disadvantages, providing  
11 kernel-mode audio processing modules.

## 12 13 **SUMMARY OF THE INVENTION**

14       Kernel-mode audio processing modules are described herein.

15       According to one aspect, multiple audio processing modules or filters are  
16 combined to form a module or filter graph. The graph is implemented in kernel-  
17 mode, reducing latency and jitter when handling audio data (e.g., MIDI data) by  
18 avoiding transfers of the audio data to user-mode applications for processing. A  
19 variety of different audio processing modules can be used to provide various  
20 pieces of functionality when processing audio data.

21       According to another aspect, a Feeder In filter is included to convert audio  
22 data received from a hardware driver into a data structure including a data portion  
23 that can include one of audio data, a pointer to a chain of additional data structures  
24 that include the audio data, and a pointer to a data buffer.

1 According to another aspect, a Feeder Out filter is included to convert, to a  
2 hardware driver-specific format, audio data received as part of a data structure  
3 including a data portion that can include one of audio data, a pointer to a chain of  
4 additional data structures that include the audio data, and a pointer to a data buffer.

5 According to another aspect, a Channel Group Mute filter is included to  
6 delete channel groups. Data packets corresponding to channel groups which  
7 match a filter parameter are forwarded to an allocator module for re-allocation of  
8 the memory space used by the data packets.

9 According to another aspect, a Channel Group Solo filter is included to  
10 delete all channel groups except for selected channel groups. Data packets  
11 corresponding to channel groups which do not match a filter parameter are  
12 forwarded to an allocator module for re-allocation of the memory space used by  
13 the data packets.

14 According to another aspect, a Channel Group Route filter is included to  
15 route groups of channels. The channel group identifiers for data packets  
16 corresponding to channel groups which match a filter parameter are changed to a  
17 new channel group.

18 According to another aspect, a Channel Group Map filter is included to  
19 alter channel group identifiers for multiple channel groups. The channel group  
20 identifiers for data packets corresponding to multiple source channel groups which  
21 match a filter parameter are changed to one or more different destination groups.

22 According to another aspect, a Channel Map filter to change any one or  
23 more of multiple channels to any one or more of the channels. Channels for data  
24 packets corresponding to multiple channels which match a filter parameter are  
25 changed to one or more different new channels. Additional data packets are

1 generated as necessary in the event of multiple new channels (a one to many  
2 mapping of channels).

3 According to another aspect, a Message Filter is included to delete selected  
4 message types. Data packets corresponding to message types which match a filter  
5 parameter are forwarded to an allocator module for re-allocation of the memory  
6 space used by the data packets.

7 According to another aspect, a Note Map Curve filter is included to alter  
8 note values on an individual basis. An input note to output note mapping table is  
9 used to identify, for each received data packet, what the input note is to be  
10 changed to (if anything).

11 According to another aspect, a Velocity Map Curve filter is included to  
12 alter velocity values on an individual basis. An input velocity to output velocity  
13 mapping table is used to identify, for each received data packet, what the input  
14 velocity is to be changed to (if anything).

15 According to another aspect, a Note and Velocity Map Curve filter is  
16 included to allow combined note and velocity alterations based on both the input  
17 note and velocity values – two degrees of freedom, leading to much more  
18 expressive translations. A table mapping input note and velocity combinations to  
19 output note and velocity combinations is used to identify, for each received data  
20 packet, what the input note and velocity are to be changed to (if anything).  
21 Alternatively, rather than changing the input note and velocity values, the Note  
22 and Velocity Map Curve filter may generate a new data structure that includes the  
23 new note and velocity values (from the table), and then forward both on to the next  
24 module in the graph.  
25

1 According to another aspect, a Time Palette filter is included to alter  
2 presentation times corresponding to the audio data. Presentation times can be  
3 quantized (e.g., snapped to a closest one of a set of presentation times) or anti-  
4 quantized (e.g., moved away from a set of presentation times). The presentation  
5 times can also be altered to generate a swing beat.

6 According to another aspect, a Variable Detune filter is included to alter the  
7 pitch of music by a variable offset value. The pitch of audio data corresponding to  
8 received data packets is altered by an amount that varies over time.

9 According to another aspect, an Echo filter is included to generate an echo  
10 for notes of the audio data. Additional data packets are generated that duplicate at  
11 least part of a received data packet, but increase the presentation time and decrease  
12 the velocity to generate an echo. The note values of the additional data packets  
13 may also be altered (e.g., for a spiraling up or spiraling down echo).

14 According to another aspect, a Profile System Performance filter is  
15 included to monitor and record system performance. System performance is  
16 monitored (e.g., a difference between presentation time for a data packet and the  
17 reference clock time just prior to rendering) and recorded for subsequent retrieval.

## 18 19 **BRIEF DESCRIPTION OF THE DRAWINGS**

20 The present invention is illustrated by way of example and not limitation in  
21 the figures of the accompanying drawings. The same numbers are used  
22 throughout the figures to reference like components and/or features.

23 Fig. 1 is a block diagram illustrating an exemplary system for manipulating  
24 and rendering audio data.  
25

1        Fig. 2 shows a general example of a computer that can be used in  
2 accordance with certain embodiments of the invention.

3        Fig. 3 is a block diagram illustrating an exemplary MIDI processing  
4 architecture in accordance with certain embodiments of the invention.

5        Fig. 4 is a block diagram illustrating an exemplary transform module graph  
6 module in accordance with certain embodiments of the invention.

7        Fig. 5 is a block diagram illustrating an exemplary MIDI message.

8        Fig. 6 is a block diagram illustrating an exemplary MIDI data packet in  
9 accordance with certain embodiments of the invention.

10       Fig. 7 is a block diagram illustrating an exemplary buffer for  
11 communicating MIDI data between a non-legacy application and a MIDI  
12 transform module graph module in accordance with certain embodiments of the  
13 invention.

14       Fig. 8 is a block diagram illustrating an exemplary buffer for  
15 communicating MIDI data between a legacy application and a MIDI transform  
16 module graph module in accordance with certain embodiments of the invention.

17       Fig. 9 is a block diagram illustrating an exemplary MIDI transform module  
18 graph such as may be used in accordance with certain embodiments of the  
19 invention.

20       Fig. 10 is a block diagram illustrating another exemplary MIDI transform  
21 module graph such as may be used in accordance with certain embodiments of the  
22 invention.

23       Fig. 11 is a flowchart illustrating an exemplary process for the operation of  
24 a module in a MIDI transform module graph in accordance with certain  
25 embodiments of the invention.



1 Fig. 12 is a flowchart illustrating an exemplary process for the operation of  
2 a graph builder in accordance with certain embodiments of the invention.

3 Fig. 13 is a block diagram illustrating an exemplary set of additional  
4 transform modules that can be made added to a module graph in accordance with  
5 certain embodiments of the invention.

6 Fig. 14 illustrates an exemplary matrix for use in a Channel Map module in  
7 accordance with certain embodiments of the invention.

## 8 9 **DETAILED DESCRIPTION**

### 10 **General Environment**

11 Fig. 1 is a block diagram illustrating an exemplary system for manipulating  
12 and rendering audio data. One type of audio data is defined by the MIDI (Musical  
13 Instrument Digital Interface) standard, including both accepted versions of the  
14 standard and proposed versions for future adoption. Although various  
15 embodiments of the invention are discussed herein with reference to the MIDI  
16 standard, other audio data standards can alternatively be used. In addition, other  
17 types of audio control information can also be passed, such as volume change  
18 messages, audio pan change messages (e.g., changing the manner in which the  
19 source of sound appears to move from two or more speakers), a coordinate change  
20 on a 3D sound buffer, messages for synchronized start of multiple devices, or any  
21 other parameter of how the audio is being processed.

22 Audio system 100 includes a computing device 102 and an audio output  
23 device 104. Computing device 102 represents any of a wide variety of computing  
24 devices, such as conventional desktop computers, gaming devices, Internet  
25 appliances, etc. Audio output device 104 is a device that renders audio data,

1 producing audible sounds based on signals received from computing device 102.  
2 Audio output device 104 can be separate from computing device 102 (but coupled  
3 to device 102 via a wired or wireless connection), or alternatively incorporated  
4 into computing device 102. Audio output device 104 can be any of a wide variety  
5 of audible sound-producing devices, such as an internal personal computer  
6 speaker, one or more external speakers, etc.

7 Computing device 102 receives MIDI data for processing, which can  
8 include manipulating the MIDI data, playing (rendering) the MIDI data, storing  
9 the MIDI data, transporting the MIDI data to another device via a network, etc.  
10 MIDI data can be received from a variety of devices, examples of which are  
11 illustrated in Fig. 1. MIDI data can be received from a keyboard 106 or other  
12 musical instruments 108 (e.g., drum machine, synthesizer, etc.), another audio  
13 device(s) 110 (e.g., amplifier, receiver, etc.), a local (either fixed or removable)  
14 storage device 112, a remote (either fixed or removable) storage device 114,  
15 another device 116 via a network (such as a local area network or the Internet),  
16 etc. Some of these MIDI data sources can generate MIDI data (e.g., keyboard  
17 106, audio device 110, or device 116 (e.g., coming via a network)), while other  
18 sources (e.g., storage device 112 or 114, or device 116) may simply be able to  
19 transmit MIDI data that has been generated elsewhere.

20 In addition to being sources of MIDI data, devices 106 – 116 may also be  
21 destinations for MIDI data. Some of the sources (e.g., keyboard 106, instruments  
22 108, device 116, etc.) may be able to render (and possibly store) the audio data,  
23 while other sources (e.g., storage devices 112 and 114) may only be able store the  
24 MIDI data.

1       The MIDI standard describes a technique for representing a musical piece  
2 as a sequence of discrete notes and other events (e.g., such as might be performed  
3 by an instrumentalist). These notes and events (the MIDI data) are communicated  
4 in messages that are typically two or three bytes in length. These messages are  
5 commonly classified as Channel Voice Messages, Channel Mode Messages, or  
6 System Messages. Channel Voice Messages carry musical performance data  
7 (corresponding to a specific channel), Channel Mode Messages affect the way a  
8 receiving instrument will respond to the Channel Voice Messages, and System  
9 Messages are control messages intended for all receivers in the system and are not  
10 channel-specific. Examples of such messages include note on and note off  
11 messages identifying particular notes to be turned on or off, aftertouch messages  
12 (e.g., indicating how long a keyboard key has been held down after being pressed),  
13 pitch wheel messages indicating how a pitch wheel has been adjusted, etc.  
14 Additional information regarding the MIDI standard is available from the MIDI  
15 Manufacturers Association of La Habra, California.

16       In the discussion herein, embodiments of the invention are described in the  
17 general context of computer-executable instructions, such as program modules,  
18 being executed by one or more conventional personal computers. Generally,  
19 program modules include routines, programs, objects, components, data structures,  
20 etc. that perform particular tasks or implement particular abstract data types.  
21 Moreover, those skilled in the art will appreciate that various embodiments of the  
22 invention may be practiced with other computer system configurations, including  
23 hand-held devices, gaming consoles, Internet appliances, multiprocessor systems,  
24 microprocessor-based or programmable consumer electronics, network PCs,  
25 minicomputers, mainframe computers, and the like. In a distributed computer

1 environment, program modules may be located in both local and remote memory  
2 storage devices.

3 Alternatively, embodiments of the invention can be implemented in  
4 hardware or a combination of hardware, software, and/or firmware. For example,  
5 at least part of the invention can be implemented in one or more application  
6 specific integrated circuits (ASICs) or programmable logic devices (PLDs).

7 Fig. 2 shows a general example of a computer 142 that can be used in  
8 accordance with certain embodiments of the invention. Computer 142 is shown as  
9 an example of a computer that can perform the functions of computing device 102  
10 of Fig. 1.

11 Computer 142 includes one or more processors or processing units 144, a  
12 system memory 146, and a bus 148 that couples various system components  
13 including the system memory 146 to processors 144. The bus 148 represents one  
14 or more of any of several types of bus structures, including a memory bus or  
15 memory controller, a peripheral bus, an accelerated graphics port, and a processor  
16 or local bus using any of a variety of bus architectures. The system memory  
17 includes read only memory (ROM) 150 and random access memory (RAM) 152.  
18 A basic input/output system (BIOS) 154, containing the basic routines that help to  
19 transfer information between elements within computer 142, such as during start-  
20 up, is stored in ROM 150.

21 Computer 142 further includes a hard disk drive 156 for reading from and  
22 writing to a hard disk, not shown, connected to bus 148 via a hard disk driver  
23 interface 157 (e.g., a SCSI, ATA, or other type of interface); a magnetic disk drive  
24 158 for reading from and writing to a removable magnetic disk 160, connected to  
25 bus 148 via a magnetic disk drive interface 161; and an optical disk drive 162 for

1 reading from or writing to a removable optical disk 164 such as a CD ROM, DVD,  
2 or other optical media, connected to bus 148 via an optical drive interface 165.  
3 The drives and their associated computer-readable media provide nonvolatile  
4 storage of computer readable instructions, data structures, program modules and  
5 other data for computer 142. Although the exemplary environment described  
6 herein employs a hard disk, a removable magnetic disk 160 and a removable  
7 optical disk 164, it should be appreciated by those skilled in the art that other types  
8 of computer readable media which can store data that is accessible by a computer,  
9 such as magnetic cassettes, flash memory cards, digital video disks, random access  
10 memories (RAMs) read only memories (ROM), and the like, may also be used in  
11 the exemplary operating environment.

12 A number of program modules may be stored on the hard disk, magnetic  
13 disk 160, optical disk 164, ROM 150, or RAM 152, including an operating system  
14 170, one or more application programs 172, other program modules 174, and  
15 program data 176. A user may enter commands and information into computer  
16 142 through input devices such as keyboard 178 and pointing device 180. Other  
17 input devices (not shown) may include a microphone, joystick, game pad, satellite  
18 dish, scanner, or the like. These and other input devices are connected to the  
19 processing unit 144 through an interface 168 that is coupled to the system bus. A  
20 monitor 184 or other type of display device is also connected to the system bus  
21 148 via an interface, such as a video adapter 186. In addition to the monitor,  
22 personal computers typically include other peripheral output devices (not shown)  
23 such as speakers and printers.

24 Computer 142 optionally operates in a networked environment using  
25 logical connections to one or more remote computers, such as a remote computer

1 188. The remote computer 188 may be another personal computer, a server, a  
2 router, a network PC, a peer device or other common network node, and typically  
3 includes many or all of the elements described above relative to computer 142,  
4 although only a memory storage device 190 has been illustrated in Fig. 2. The  
5 logical connections depicted in Fig. 2 include a local area network (LAN) 192 and  
6 a wide area network (WAN) 194. Such networking environments are  
7 commonplace in offices, enterprise-wide computer networks, intranets, and the  
8 Internet. In the described embodiment of the invention, remote computer 188  
9 executes an Internet Web browser program (which may optionally be integrated  
10 into the operating system 170) such as the "Internet Explorer" Web browser  
11 manufactured and distributed by Microsoft Corporation of Redmond, Washington.

12 When used in a LAN networking environment, computer 142 is connected  
13 to the local network 192 through a network interface or adapter 196. When used  
14 in a WAN networking environment, computer 142 typically includes a modem 198  
15 or other component for establishing communications over the wide area network  
16 194, such as the Internet. The modem 198, which may be internal or external, is  
17 connected to the system bus 148 via an interface (e.g., a serial port interface 168).  
18 In a networked environment, program modules depicted relative to the personal  
19 computer 142, or portions thereof, may be stored in the remote memory storage  
20 device. It is to be appreciated that the network connections shown are exemplary  
21 and other means of establishing a communications link between the computers  
22 may be used.

23 Computer 142 also optionally includes one or more broadcast tuners 200.  
24 Broadcast tuner 200 receives broadcast signals either directly (e.g., analog or  
25

1 digital cable transmissions fed directly into tuner 200) or via a reception device  
2 (e.g., via antenna 110 or satellite dish 114 of Fig. 1).

3       Generally, the data processors of computer 142 are programmed by means  
4 of instructions stored at different times in the various computer-readable storage  
5 media of the computer. Programs and operating systems are typically distributed,  
6 for example, on floppy disks or CD-ROMs. From there, they are installed or  
7 loaded into the secondary memory of a computer. At execution, they are loaded at  
8 least partially into the computer's primary electronic memory. The invention  
9 described herein includes these and other various types of computer-readable  
10 storage media when such media contain instructions or programs for implementing  
11 the steps described below in conjunction with a microprocessor or other data  
12 processor. The invention also includes the computer itself when programmed  
13 according to the methods and techniques described below. Furthermore, certain  
14 sub-components of the computer may be programmed to perform the functions  
15 and steps described below. The invention includes such sub-components when  
16 they are programmed as described. In addition, the invention described herein  
17 includes data structures, described below, as embodied on various types of  
18 memory media.

19       For purposes of illustration, programs and other executable program  
20 components such as the operating system are illustrated herein as discrete blocks,  
21 although it is recognized that such programs and components reside at various  
22 times in different storage components of the computer, and are executed by the  
23 data processor(s) of the computer.

## **Kernel-Mode Processing**

Fig. 3 is a block diagram illustrating an exemplary MIDI processing architecture in accordance with certain embodiments of the invention. The architecture 308 includes application(s) 310, graph builder 312, a MIDI transform module graph 314, and hardware devices 316 and 318. Hardware devices 316 and 318 are intended to represent any of a wide variety of MIDI data input and/or output devices, such as any of devices 104 – 116 of Fig. 1. Hardware devices 316 and 318 are implemented in hardware level 320 of architecture 308.

Hardware devices 316 and 318 communicate with MIDI transform module graph 314, passing input data to modules in graph 314 and receiving data from modules in graph 314. Hardware devices 316 and 318 communicate with modules in MIDI transform module graph 314 via hardware (HW) drivers 322 and 324, respectively. A portion of each of hardware drivers 322 and 324 is implemented as a module in graph 314 (these portions are often referred to as "miniport streams"), and a portion is implemented in software external to graph 314 (often referred to as "miniport drivers"). For input of MIDI data from a hardware device 316 (or 318), the hardware driver 322 (or 324) reads the data off of the hardware device 316 (or 318) and puts the data in a form expected by the modules in graph 314. For output of MIDI data to a hardware device 316 (or 318), the hardware driver receives the data and writes this data to the hardware directly.

An additional "feeder" module may also be included that is situated between the miniport stream and the rest of the graph 314. Such feeder modules are particularly useful in situations where the miniport driver is not aware of the graph 314 or the data formats and protocols used within graph 314. In such situations, the feeder module operates to convert formats between the hardware



1 (and hardware driver) specific format and the format supported by graph 314.  
2 Essentially, for older miniport drivers whose miniport streams don't communicate  
3 in the format supported by graph 314, the FeederIn and FeederOut modules  
4 function as their liaison into that graph.

5 MIDI transform module graph 314 includes multiple ( $n$ ) modules 326 (also  
6 referred to as filters or MXFs (MIDI transform filters)) that can be coupled  
7 together. Different source to destination paths (e.g., hardware device to hardware  
8 device, hardware device to application, application to hardware device, etc.) can  
9 exist within graph 314, using different modules 326 or sharing modules 326. Each  
10 module 326 performs a particular function in processing MIDI data. Examples of  
11 modules 326 include a sequencer to control the output of MIDI data to hardware  
12 device 316 or 318 for playback, a packer module to package MIDI data for output  
13 to application 310, etc. The operation of modules 326 is discussed in further detail  
14 below.

15 Modern operating systems (e.g., those in the Microsoft Windows® family  
16 of operating systems) typically include multiple privilege levels, often referred to  
17 as user and kernel modes of operation (also called "ring 3" and "ring 0"). Kernel-  
18 mode is usually associated with and reserved for portions of the operating system.  
19 Kernel-mode (or "ring 0") components run in a reserved address space, which is  
20 protected from user-mode components. User-mode (or "ring 3") components have  
21 their own respective address spaces, and can make calls to kernel-mode  
22 components using special procedures that require so-called "ring transitions" from  
23 one privilege level to another. A ring transition involves a change in execution  
24 context, which involves not only a change in address spaces, but also a transition  
25 to a new processor state (including register values, stacks, privilege mode, etc).

1 As discussed above, such ring transitions can result in considerable latency and an  
2 unpredictable amount of time.

3 MIDI transform module graph 314 is implemented in kernel-mode of  
4 software level 328. Modules 326 are all implemented in kernel-mode, so no ring  
5 transitions are required during the processing of MIDI data. Modules 326 are  
6 implemented at a deferred procedure call (DPC) level, such as  
7 DISPATCH\_LEVEL. By implementing modules 326 at a higher priority level  
8 than other user-mode software components, the modules 326 will have priority  
9 over the user-mode components, thereby reducing delays in executing modules  
10 326 and thus reducing latency and unpredictability in the transmitting and  
11 processing of MIDI data.

12 In the illustrated example, modules 326 are implemented using Win32®  
13 Driver Model (WDM) Kernel Streaming filters, thereby reducing the amount of  
14 overhead necessary in communicating between modules 326. A low-overhead  
15 interface is used by modules 326 to communicate with one another, rather than  
16 higher-overhead I/O Request Packets (IRPs), and is described in more detail  
17 below. Additional information regarding the WDM Kernel Streaming architecture  
18 is available from Microsoft Corporation of Redmond, Washington.

19 Software level 328 also includes application(s) 310 implemented in user-  
20 mode, and graph builder 312 implemented in kernel-mode. Any number of  
21 applications 310 can interface with graph 314 (concurrently, in the event of a  
22 multi-tasking operating system). Application 310 represents any of a wide variety  
23 of applications that may use MIDI data. Examples of such applications include  
24 games, reference materials (e.g., dictionaries or encyclopedias) and audio  
25 programs (e.g., audio player, audio mixer, etc.).

1 In the illustrated example, graph builder 312 is responsible for generating a  
2 particular graph 314. MIDI transform module graph 314 can vary depending on  
3 what MIDI processing is desired. For example, a pitch modification module 326  
4 would be included in graph 314 if pitch modification is desired, but otherwise  
5 would not be included. MIDI transform module graph 314 has multiple different  
6 modules available to it, although only selected modules may be incorporated into  
7 graph 314 at any particular time. In the illustrated example, MIDI transform  
8 module graph 314 can include multiple modules 326 that do not have connections  
9 to other modules 326 – they simply do not operate on received MIDI data.  
10 Alternatively, only modules that operate on received MIDI data may be included  
11 in graph 314, with graph builder 312 accessing a module library 330 to copy  
12 modules into graph 314 when needed.

13 In one implementation, graph builder 312 accesses one or more locations to  
14 identify which modules are available to it. By way of example, a system registry  
15 may identify the modules or an index associated with module library 330 may  
16 identify the modules. Whenever a new module is added to the system, an  
17 identification of the module is added to these one or more locations. The  
18 identification may also include a descriptor, usable by graph builder 312 and/or an  
19 application 310, to identify the type of functionality provided by the module.

20 Graph builder 312 communicates with the individual modules 326 to  
21 configure graph 314 to carry out the desired MIDI processing functionality, as  
22 indicated to graph builder 312 by application 310. Although illustrated as a  
23 separate application that is accessed by other user-mode applications (e.g.,  
24 application 310), graph builder 312 may alternatively be implemented as part of  
25

1 another application (e.g., part of application 310), or may be implemented as a  
2 separate application or system process in user-mode.

3 Application 310 can determine what functionality should be included in  
4 MIDI transform module graph 314 (and thus what modules graph builder 312  
5 should include in graph 314) in any of a wide variety of manners. By way of  
6 example, application 310 may provide an interface to a user (e.g., a graphical user  
7 interface) that allows the user to identify various alterations he or she would like  
8 made to a musical piece. By way of another example, application 310 may be pre-  
9 programmed with particular functionality of what alterations should be made to a  
10 musical piece, or may access another location (e.g., a remote server computer) to  
11 obtain the information regarding what alterations should be made to the musical  
12 piece. Additionally, graph builder 312 may automatically insert certain  
13 functionality into the graph, as discussed in more detail below.

14 Graph builder 312 can change the connections in MIDI transform module  
15 graph 314 during operation of the graph. In one implementation, graph builder  
16 312 pauses or stops operation of graph 314 temporarily in order to make the  
17 necessary changes, and then resumes operation of the graph. Alternatively, graph  
18 builder 312 may change connections in the graph without stopping its operation.  
19 Graph builder 312 and the manner in which it manages graph 314 are discussed in  
20 further detail below.

21 MIDI transform module graphs are thus readily extensible. Graph builder  
22 312 can re-arrange the graph in any of a wide variety of manners to accommodate  
23 the desires of an application 310. New modules can be incorporated into a graph  
24 to process MIDI data, modules can be removed from the graph so they no longer  
25

1 process MIDI data, connections between modules can be modified so that modules  
2 pass MIDI data to different modules, etc.

3 Communication between applications 310 and MIDI transform module  
4 graph 314 transitions between different rings, so some latency and temporal  
5 unpredictability may be experienced. In one implementation, communication  
6 between applications 310 (or graph builder 312) and a module 326 is performed  
7 using conventional IRPs. However, the processing of the MIDI data is being  
8 carried out in kernel-mode, so such latency and/or temporal unpredictability does  
9 not adversely affect the processing of the MIDI data.

10 Fig. 4 is a block diagram illustrating an exemplary module 326 in  
11 accordance with certain embodiments of the invention. In the illustrated example,  
12 each module 326 in graph 314 includes a processing portion 332 in which the  
13 operation of the module 326 is carried out (and which varies by module). Each  
14 module 326 also includes four interfaces: SetState 333, PutMessage 334,  
15 ConnectOutput 335, and DisconnectOutput 336.

16 The SetState interface 333 allows the state of a module 326 to be set (e.g.,  
17 by an application 310 or graph builder 312). In one implementation, valid states  
18 include run, acquire, pause, and stop. The run state indicates that the module is to  
19 run and perform its particular function. The acquire and pause states are  
20 transitional states that can be used to assist in transitioning between the run and  
21 stop states. The stop state indicates that the module is to stop running (it won't  
22 accept any inputs or provide any outputs). When the SetState interface 333 is  
23 called, one of the four valid states is included as a parameter by the calling  
24 component.

1       The PutMessage interface 334 allows MIDI data to be input to a module  
2 326. When the PutMessage interface 334 is called by another module, a pointer to  
3 the MIDI data being passed (e.g., a data packet, as discussed in more detail below)  
4 is included as a parameter, allowing the pointer to the MIDI data to be forwarded  
5 to processing portion 332 for processing of the MIDI data. The PutMessage  
6 interface 334 is called by another module 326, after it has finished processing the  
7 MIDI data it received, and which passes the processed MIDI data to the next  
8 module in the graph 314. After processing portion 332 finishes processing the  
9 MIDI data, the PutMessage interface on the next module in the graph is called by  
10 processing portion 332 to transfer the processed MIDI data to the connected  
11 module 326 (the next module in the graph, as discussed below).

12       The ConnectOutput interface 335 allows a module 326 to be programmed  
13 with the connected module (the next module in the graph). The ConnectOutput  
14 interface is called by graph builder 312 to identify to the module where the output  
15 of the module should be sent. When the ConnectOutput interface 335 is called, an  
16 identifier (e.g., pointer to) the next module in the graph is included as a parameter  
17 by the calling component. The default connected output is the allocator (discussed  
18 in more detail below). In one implementation (called a “splitter” module), a  
19 module 326 can be programmed with multiple connected modules (e.g., by  
20 programming the module 326 with the PutMessage interfaces of each of the  
21 multiple connected modules), allowing outputs to multiple “next” modules in the  
22 graph. Conversely, multiple modules can point at a single “next” output module  
23 (e.g., multiple modules may be programmed with the PutMessage interface of the  
24 same “next” module).

1       The DisconnectOutput interface 336 allows a module 326 to be  
2       disconnected from whatever module it was previously connected to (via the  
3       ConnectOutput interface). The DisconnectOutput interface 336 is called by graph  
4       builder 312 to have the module 326 reset to a default connected output (the  
5       allocator). When the DisconnectOutput interface 336 is called, an identifier (e.g.,  
6       pointer to) the module being disconnected from is included as a parameter by the  
7       calling component. In one implementation, calling the ConnectOutput interface  
8       335 or DisconnectOutput interface 336 with a parameter of NULL also  
9       disconnects the “next” reference. Alternatively, the DisconnectOutput interface  
10      336 may not be included (e.g., disconnecting the module can be accomplished by  
11      calling ConnectOutput 335 with a NULL parameter, or with an identification of  
12      the allocator module as the next module).

13      Additional interfaces 337 may also be included on certain modules,  
14      depending on the functions performed by the module. Two such additional  
15      interfaces 337 are illustrated in Fig. 4: a SetParameters interface 338 and a  
16      GetParameters interface 339. The SetParameters interface 338 allows a module  
17      326 to receive various operational parameters set (e.g., from applications 310 or  
18      graph builder 312), which are maintained as parameters 340. For example, a  
19      module 326 that is to alter the pitch of a particular note(s) can be programmed, via  
20      the SetParameters interface 338, with which note is to be altered and/or how much  
21      the pitch is to be altered.

22      The GetParameters interface 339 allows coefficients (e.g., operational  
23      parameters maintained as parameters 340) previously sent to the module, or any  
24      other information the module may have been storing in a data section 341 (such as  
25      MIDI jitter performance profiling data, number of events left in the allocator’s free

1 memory pool, how much memory is currently allocated by the allocator, how  
2 many messages have been enqueued by a sequencer module, a breakdown by  
3 channel and/or channel group of what messages have been enqueued by the  
4 sequencer module, etc), to be retrieved. The GetParameters interface 339 and  
5 SetParameters interface 338 are typically called by graph builder 312, although  
6 other applications 310 or modules in graph 314 could alternatively call them.

7       Returning to Fig. 3, one particular module that is included in MIDI  
8 transform module graph 314 is referred to as the allocator. The allocator module  
9 is responsible for obtaining memory from the memory manager (not shown) of the  
10 computing device and making portions of the obtained memory available for  
11 MIDI data. The allocator module makes a pool of memory available for allocation  
12 to other modules in graph 314 as needed. The allocator module is called by  
13 another module 326 when MIDI data is received into the graph 314 (e.g., from  
14 hardware device 316 or 318, or application 310). The allocator module is also  
15 called when MIDI data is transferred out of the graph 314 (e.g., to hardware  
16 device 316 or 318, or application 310) so that memory that was being used by the  
17 MIDI data can be reclaimed and re-allocated for use by other MIDI data.

18       The allocator includes the interfaces discussed above, as well as additional  
19 interfaces that differ from the other modules 326. In the illustrated example, the  
20 allocator includes four additional interfaces: GetMessage, GetBufferSize,  
21 GetBuffer, and PutBuffer.

22       The GetMessage interface is called by another module 326 to obtain a data  
23 structure into which MIDI data can be input. The modules 326 communicate  
24 MIDI data to one another using a structure referred to as a data packet or event.  
25 Calling the GetMessage interface causes the allocator to return to the calling



1 module a pointer to such a data packet in which the calling module can store MIDI  
2 data.

3 The PutMessage interface for the allocator takes a data structure and returns  
4 it to the free pool of packets that it maintains. This consists of its “processing.”  
5 The allocator is the original source and the ultimate destination of all event data  
6 structures of this type.

7 MIDI data is typically received in two or three byte messages. However,  
8 situations can arise where larger portions of MIDI data are received, referred to as  
9 System Exclusive, or SysEx messages. In such situations, the allocator allocates a  
10 larger buffer for the MIDI data, such as 60 bytes or 4096 bytes. The  
11 GetBufferSize interface is called by a module 326, and the allocator responds with  
12 the size of the buffer that is (or will be) allocated for the portion of data. In one  
13 implementation, the allocator always allocates buffers of the same size, so the  
14 response by the allocator is always the same.

15 The GetBuffer interface is called by a module 326 and the allocator  
16 responds by passing, to the module, a pointer to the buffer that can be used by the  
17 module for the portion of MIDI data.

18 The PutBuffer interface is called by a module 326 to return the memory  
19 space for the buffer to the allocator for re-allocation (the PutMessage interface  
20 described above will call PutBuffer in turn, to return the memory space to the  
21 allocator, if this hasn’t been done already). When calling the PutBuffer interface,  
22 the calling module includes, as a parameter, a pointer to the buffer being returned  
23 to the allocator.

24 Situations can also arise where the amount of memory that is allocated by  
25 the allocator for a buffer is smaller than the portion of MIDI data that is to be

1 received. In this situation, multiple buffers are requested from the allocator and  
2 are "chained" together (e.g., a pointer in a data packet corresponding to each  
3 identifies the starting point of the next buffer). An indication may also be made in  
4 the corresponding data packet that identifies whether a particular buffer stores the  
5 entire portion of MIDI data or only a sub-portion of the MIDI data.

6 Many modern processors and operating systems support virtual memory.  
7 Virtual memory allows the operating system to allocate more memory to  
8 application processes than is physically available in the computing device. Data  
9 can then be swapped between physical memory (e.g., RAM) and another storage  
10 device (e.g., a hard disk drive), a process referred to as paging. The use of virtual  
11 memory gives the appearance of more physical memory being available in the  
12 computing device than is actually available. The tradeoff, however, is that  
13 swapping data from a disk drive to memory typically takes significantly longer  
14 than simply retrieving the data directly from memory.

15 In one implementation, the allocator obtains non-pageable portions of  
16 memory from the memory manager. That is, the memory that is obtained by the  
17 allocator refers to a portion of physical memory that will not be swapped to disk.  
18 Thus, processing of MIDI data will not be adversely affected by delays in  
19 swapping data between memory and a disk.

20 In one implementation, each module 326, when added to graph 314, is  
21 passed an identifier (e.g., pointer to) the allocator module as well as a clock. The  
22 allocator module is used, as described above, to allow memory for MIDI data to be  
23 obtained and released. The clock is a common reference clock that is used by all  
24 of the modules 326 to maintain synchronization with one another. The manner in  
25 which the clock is used can vary, depending on the function performed by the

1 modules. For example, a module may generate a time stamp, based on the clock,  
2 indicating when the MIDI data was received by the module, or may access a  
3 presentation time for the data indicating when it is to be played back.

4 Alternatively, some modules may not need, and thus need not include,  
5 pointers to the reference clock and/or the allocator module (however, in  
6 implementations where the default output destination for each module is an  
7 allocator module, then each module needs a pointer to the allocator in order to  
8 properly initialize). For example, if a module will carry out its functionality  
9 without regard for what the current reference time is, then a pointer to the  
10 reference clock is not necessary.

11 Fig. 5 is a block diagram illustrating an exemplary MIDI message 345.  
12 MIDI message 345 includes a status portion 346 and a data portion 347. Status  
13 portion 346 is one byte, while data portion 347 is either one or two bytes. The size  
14 of data portion 347 is encoded in the status portion 346 (either directly, or  
15 inherently based on some other value (such as the type of command)). The MIDI  
16 data is received from and passed to hardware devices 316 and 318 of Fig. 3, and  
17 possibly application 310, as messages 345. Typically each message 345 identifies  
18 a single command (e.g., note on, note off, change volume, pitch bend, etc.). The  
19 audio data included in data portion 347 will vary depending on the message type.

20 Fig. 6 is a block diagram illustrating an exemplary MIDI data packet 350 in  
21 accordance with certain embodiments of the invention. MIDI data (or references,  
22 such as pointers, thereto) is communicated among modules 326 in MIDI transform  
23 module graph 314 of Fig. 3 as data packets 350, also referred to as events. When a  
24 MIDI message 345 of Fig. 5 is received into graph 314, the receiving module 326  
25 generates a data packet 350 that incorporates the message.

1 Data packet 350 includes a reserved portion 352 (e.g., one byte), a structure  
2 byte count portion 354 (e.g., one byte), an event byte count portion 356 (e.g. two  
3 bytes), a channel group portion 358 (e.g., two bytes), a flags portion 360 (e.g. two  
4 bytes), a presentation time portion 362 (e.g., eight bytes), a byte position 364 (e.g.,  
5 eight bytes), a next event portion 366 (e.g. four bytes), and a data portion 368  
6 (e.g., four bytes). Reserved portion 352 is reserved for future use. Structure byte  
7 count portion 354 identifies the size of the message 350.

8 Event byte count portion 356 identifies the number of data bytes that are  
9 referred to in data portion 368. The number of data bytes could be the number  
10 actually stored in data portion 368 (e.g., two or three, depending on the type of  
11 MIDI data), or alternatively the number of bytes pointed to by a pointer in data  
12 portion 368, (e.g., if the number of data bytes is greater than the size of a pointer).  
13 If the event is a package event (pointing to a chain of events, as discussed in more  
14 detail below), then the portion 356 has no value. Alternatively, portion 356 could  
15 be set to the value of event byte count portion 356 of the first regular event in its  
16 chain, or to the byte count of the entire long message. If event portion 356 is not  
17 set to the byte count of the entire long message, then data could still be flowing  
18 into the last message structure of the package event while the initial data is already  
19 being processed elsewhere.

20 Channel group portion 358 identifies which of multiple channel groups the  
21 data identified in data portion 368 corresponds to. The MIDI standard supports  
22 sixteen different channels, allowing essentially sixteen different instruments or  
23 "voices" to be processed and/or played concurrently for a musical piece. Use of  
24 channel groups allows the number of channels to be expanded beyond sixteen.  
25 Each channel group can refer to any one of sixteen channels (as encoded in status

1 byte 346 of message 345 of Fig. 5). In one implementation, channel group portion  
2 358 is a 2-byte value, allowing up to 65,536 (64k) different channel groups to be  
3 identified (as each channel group can have up to sixteen channels, this allows a  
4 total of 1,048,576 (1Meg) different channels).

5 Flags portion 360 identifies various flags that can be set regarding the MIDI  
6 data corresponding to data packet 350. In one implementation, zero or more of  
7 multiple different flags can be set: an Event In Use (EIU) flag, an Event  
8 Incomplete (EI) flag, one or more MIDI Parse State flags (MPS), or a Package  
9 Event (PE) flag. The Event In Use flag should always be on (set) when an event is  
10 traveling through the system; when it is in the free pool this bit should be cleared.  
11 This is used to prevent memory corruption. The Event Incomplete flag is set if the  
12 event continues beyond the buffer pointed to by data portion 368, or if the  
13 message is a System Exclusive (SysEx) message. The MIDI Parse State flags are  
14 used by a capture sink module (or other module parsing an unparsed stream of  
15 MIDI data) in order to keep track of the state of the unparsed stream of MIDI data.  
16 As the capture sink module successfully parses the MIDI data into a complete  
17 message, these two bits should be cleared. In one implementation these flags have  
18 been removed from the public flags field.

19 The Package Event flag is set if data packet 350 points to a chain of other  
20 packets 350 that should be dealt with atomically. By way of example, if a portion  
21 of MIDI data is being processed that is large enough to require a chain of data  
22 packets 350, then this packet chain should be passed around atomically (e.g., not  
23 separated so that a module receives only a portion of the chain). Setting the  
24 Package Event flag identifies data field 374 as pointing to a chain of multiple  
25 additional packets 350.

1       Presentation time portion 362 specifies the presentation time for the data  
2 corresponding to data packet 350 (i.e., for an event). The presentation of an event  
3 depends on the type of event: note on events are presented by rendering the  
4 identified note, note off events are presented by ceasing rendering of the identified  
5 note, pitch bend events are presented by altering the pitch of the identified note in  
6 the identified manner, etc. A module 326 of Fig. 3, by comparing the current  
7 reference clock time to the presentation time identified in portion 362, can  
8 determine when, relative to the current time, the event should be presented to a  
9 hardware device 316 or 318. In one implementation, portion 362 identifies  
10 presentation times in 100 nanosecond (ns) units.

11       Byte position portion 364 identifies where this message (included in data  
12 portion 368) is situated in the overall stream of bytes from the application (e.g.,  
13 application 310 of Fig. 3). Because certain applications use the release of their  
14 submitted buffers as a timing mechanism, it is important to keep track of how far  
15 processing has gone in the byte order, and release buffers only up to that point  
16 (and only release those buffers back to the application after the corresponding  
17 bytes have actually been played). In this case the allocator module looks at the  
18 byte offset when a message is destroyed (returned for re-allocation), and alerts a  
19 stream object (e.g., the IRP stream object used to pass the buffer to graph 314) that  
20 a certain amount of memory can be released up to the client application.

21       Next event portion 366 identifies the next packet 350 in a chain of packets,  
22 if any. If there is no next packet, then next event portion 366 is NULL.

23       Data portion 368 can include one of three things: packet data 370 (a  
24 message 345 of Fig. 5), a pointer 372 to a chain of packets 350, or a pointer 374 to  
25 a data buffer. Which of these three things is included in data portion 368 can be

1 determined based on the value in event byte count field 356 and/or flags portion  
2 360. In the illustrated example, the size of a pointer is greater than three bytes  
3 (e.g., is 4 bytes). If the event byte count field 356 is less than or equal to the size  
4 of a pointer, then data portion 368 includes packet data 370; otherwise data portion  
5 368 includes a pointer 374 to a data buffer. However, this determination is  
6 overridden if the Package Event flag of flags portion 360 is set, which indicates  
7 that data portion 368 includes a pointer 372 to a chain of packets (regardless of the  
8 value of event byte count field 356).

9       Returning to Fig. 3, certain modules 326 may receive MIDI data from  
10 application 310 and/or send MIDI data to application 310. In the illustrated  
11 example, MIDI data can be received from and/or sent to an application 310 in  
12 different formats, depending at least in part on whether application 310 is aware of  
13 the MIDI transform module graph 314 and the format of data packets 350 (of Fig.  
14 5) used in graph 314. If application 310 is not aware of the format of data packets  
15 350 then application 310 is referred to as a "legacy" application and the MIDI data  
16 received from application 310 is converted into the format of data packets 350.  
17 Application 310, whether a legacy application or not, communicates MIDI data to  
18 (or receives MIDI data from) a module 326 in a buffer including one or more  
19 MIDI messages (or data packets 350).

20       Fig. 7 is a block diagram illustrating an exemplary buffer for  
21 communicating MIDI data between a non-legacy application and a MIDI  
22 transform module graph module in accordance with certain embodiments of the  
23 invention. A buffer 380, which can be used to store one or more packaged data  
24 packets, is illustrated including multiple packaged data packets 382 and 384. Each  
25 packaged data packet 382 and 384 includes a data packet 350 of Fig. 6 as well as

1 additional header information. This combination of data packet 350 and header  
2 information is referred to as a packaged data packet. In one implementation,  
3 packaged data packets are quadword (8-byte) aligned for alignment and speed  
4 reasons (e.g., by adding padding 394 as needed).

5 The header information for each packaged data packet includes an event  
6 byte count portion 386, a channel group portion 388, a reference time delta portion  
7 390, and a flags portion 392. The event byte count portion 386 identifies the  
8 number of bytes in the event(s) corresponding to data packet 350 (which is the  
9 same value as maintained in event portion 356 of data packet 350 of Fig. 6, unless  
10 the packet is broken up into multiple events structures.). The channel group  
11 portion 388 identifies which of multiple channel groups the event(s) corresponding  
12 to data packet 350 correspond to (which is the same value as maintained in  
13 channel group portion 358 of data packet 350).

14 The reference time delta portion 390 identifies the difference in  
15 presentation time between packaged data packet 382 (stored in presentation time  
16 portion 362 of data packet 350 of Fig. 6) and the beginning of buffer 380. The  
17 beginning time of buffer 380 can be identified as the presentation time of the first  
18 packaged data packet 382 in buffer 380, or alternatively buffer 380 may have a  
19 corresponding start time (based on the same reference clock as the presentation  
20 time of data packets 350 are based on).

21 Flags portion 392 identifies one or more flags that can be set regarding the  
22 corresponding data packet 350. In one implementation, only one flag is  
23 implemented - an Event Structured flag that is set to indicate that structured data is  
24 included in data packet 350. Structured data is expected to parse correctly from a  
25 raw MIDI data stream into complete message packets. An unstructured data



1 stream is perhaps not MIDI compliant, so it isn't grouped into MIDI messages like  
2 a structured stream is – the original groupings of bytes of unstructured data are  
3 unmodified. Whether the data is compliant (structured) or non-compliant  
4 (unstructured) is indicated by the Event Structured flag.

5 Fig. 8 is a block diagram illustrating an exemplary buffer for  
6 communicating MIDI data between a legacy application and a MIDI transform  
7 module graph module in accordance with certain embodiments of the invention.  
8 A buffer 410, which can be used to store one or more packaged events, is  
9 illustrated including multiple packaged events 412 and 414. Each packaged event  
10 412 and 414 includes a message 345 of Fig. 5 as well as additional header  
11 information. This combination of message 345 and header information is referred  
12 to as a packaged event (or packaged message). In one implementation, packaged  
13 events are quadword (8-byte) aligned for speed and alignment reasons (e.g., by  
14 adding padding 420 as needed).

15 The additional header information in each packaged event includes a time  
16 delta portion 416 and a byte count portion 418. Time delta portion 416 identifies  
17 the difference between the presentation time of the packaged event and the  
18 presentation time of the immediately preceding packaged event. These  
19 presentation times are established by the legacy application passing the MIDI data  
20 to the graph. For the first packaged event in buffer 410, time delta portion 416  
21 identifies the difference between the presentation time of the packed event and the  
22 beginning time corresponding to buffer 410. The beginning time corresponding to  
23 buffer 410 is the presentation time for the entire buffer (the first message in the  
24 buffer can have some positive offset in time and does not have to start right at the  
25 head of the buffer).

1           Byte count portion 416 identifies the number of bytes in message 345.

2           Fig. 9 is a block diagram illustrating an exemplary MIDI transform module  
3 graph 430 such as may be used in accordance with certain embodiments of the  
4 invention. In the illustrated example, keys on a keyboard can be activated and the  
5 resultant MIDI data forwarded to an application executing in user-mode as well as  
6 being immediately played back. Additionally, MIDI data can be input to graph  
7 430 from a user-mode application for playback.

8           One source of MIDI data in Fig. 9 is keyboard 432, which provides the  
9 MIDI data as a raw stream of MIDI bytes via a hardware driver including a  
10 miniport stream (in) module 434. Module 434 calls the GetMessage interface of  
11 allocator 436 for memory space (a data packet 350) into which a structured packet  
12 can be placed, and module 434 adds a timestamp to the data packet 350.  
13 Alternatively, module 434 may rely on capture sink module 438, discussed below,  
14 to generate the packets 350, in which case module 434 adds a timestamp to each  
15 byte of the raw data it receives prior to forwarding the data to capture sink module  
16 438. In the illustrated example, notes are to be played immediately upon  
17 activation of the corresponding key on keyboard 432, so the timestamp stored by  
18 module 434 as the presentation time of the data packets 350 is the current reading  
19 of the master (reference) clock.

20           Module 434 is connected to capture sink module 438, splitter module 430  
21 or packer 442 (the splitter module is optional – only inserted if, for example, the  
22 graph builder has been told to connect “kernel THRU”). Capture sink module 438  
23 is optional, and operates to generate packets 350 from a received MIDI data byte  
24 stream. If module 434 generates packets 350, then capture sink 438 is not  
25 necessary and module 434 is connected to optional splitter module 440 or packer

1 442. However, if module 434 does not generate packets 350, then module 434 is  
2 connected to capture sink module 438. After adding the timestamp, module 434  
3 calls the PutMessage interface of the module it is connected to (either capture sink  
4 module 438, splitter module 440 or packer 442), which passes the newly created  
5 message to that module.

6 The manner in which packets 350 are generated from the received raw  
7 MIDI data byte stream (regardless of whether it is performed by module 434 or  
8 capture sink module 438) is dependent on the particular type of data (e.g., the data  
9 may be included in data portion 368 (Fig. 6), a pointer may be included in data  
10 portion 368, etc.). In situations where multiple bytes of raw MIDI data are being  
11 stored in data portion 368, the timestamp of the first of the multiple bytes is used  
12 as the timestamp for the packet 350. Additionally, situations can arise where  
13 additional event structures have been obtained from allocator 436 than are actually  
14 needed (e.g., multiple bytes were not received together and multiple event  
15 structures were received for each, but they are to be grouped together in the same  
16 event structure). In such situations the additional event structures can be kept for  
17 future MIDI data, or alternatively returned to allocator 436 for re-allocation.

18 Splitter module 440 operates to duplicate received data packets 350 and  
19 forward each to a different module. In the illustrated example, splitter module 440  
20 is connected to both packer module 442 and sequencer module 444. Upon receipt  
21 of a data packet 350, splitter module 440 obtains additional memory space from  
22 allocator 436, copies the contents of the received packet into the new packet  
23 memory space, and calls the PutMessage interfaces of the modules it is connected  
24 to, which passes one data packet 350 to each of the connected modules (i.e., one  
25 data packet to packer module 442 and one data packet to sequencer module 444).

1 Splitter module 440 may optionally operate to duplicate a received data packet 350  
2 only if the received data packet corresponds to audio data matching a particular  
3 type, such as certain note(s), channel(s), and/or channel group(s).

4 Packer module 442 operates to combine one or more received packets into  
5 a buffer (such as buffer 380 of Fig. 7 or buffer 410 of Fig. 8) and forward the  
6 buffer to a user-mode application (e.g., using IRPs with a message format desired  
7 by the application). Two different packer modules can be used as packer module  
8 442, one being dedicated to legacy applications and the other being dedicated to  
9 non-legacy applications. Alternatively, a single packer module may be used and  
10 the type of buffer (e.g., buffer 380 or 410) used by packer module 442 being  
11 dependent on whether the application to receive the buffer is a legacy application.

12 Once a data packet is forwarded to the user-mode application, packer 442  
13 calls its programmed PutMessage interface (the PutMessage interface that the  
14 module packer 442 is connected to) for that packet. Packer module 442 is  
15 connected to allocator module 436, so calling its programmed PutMessage  
16 interface for a data packet returns the memory space used by the data packet to  
17 allocator 436 for re-allocation. Alternatively, packer 442 may wait to call  
18 allocator 436 for each packet in the buffer after the entire buffer is forwarded to  
19 the user-mode application.

20 Sequencer module 444 operates to control the delivery of data packets 350  
21 received from splitter module 440 to miniport stream (out) module 446 for playing  
22 on speakers 450. Sequencer module 444 does not change the data itself, but  
23 module 444 does reorder the data packets by timestamp and delay the calling of  
24 PutMessage (to forward the message on) until the appropriate time. Sequencer  
25 module 444 is connected to module 446, so calling PutMessage causes sequencer

1 module 444 to forward a data packet to module 446. Sequencer module 444  
2 compares the presentation times of received data packets 350 to the current  
3 reference time. If the presentation time is equal to or earlier than the current time  
4 then the data packet 350 is to be played back immediately and the PutMessage  
5 interface is called for the packet. However, if the presentation time is later than  
6 the current time, then the data packet 350 is queued until the presentation time is  
7 equal to the current time, at which point sequencer module 444 calls its  
8 programmed PutMessage interface for the packet. In one implementation,  
9 sequencer 444 is a high-resolution sequencer, measuring time in 100 ns units.

10 Alternatively, sequencer module 444 may attempt to forward packets to  
11 module 446 slightly in advance of their presentation time (that is, when the  
12 presentation time of the packet is within a threshold amount of time later than the  
13 current time). The amount of this threshold time would be, for example, an  
14 anticipated amount of time that is necessary for the data packet to pass through  
15 module 446 and to speakers 450 for playing, resulting in playback of the data  
16 packets at their presentation times rather than submission of the packets to module  
17 446 at their presentation times. An additional "buffer" amount of time may also be  
18 added to the anticipated amount of time to allow output module 448 (or speakers  
19 450) to have the audio messages delivered at a particular time (e.g., five seconds  
20 before the data needs to be rendered by speakers 450).

21 A module 446 could furthermore specify that it did not want the sequencer  
22 to hold back the data at all, even if data were extremely early. In this case, the  
23 HW driver "wants to do its own sequencing," so the sequencer uses a very high  
24 threshold (or alternatively a sequencer need not be inserted above this particular  
25 module 446). The module 446 is receiving events with presentation timestamps in

1 them, and it also has access to the clock (e.g., being handed a pointer to it when it  
2 was initialized), so if the module 446 wanted to synchronize that clock to its own  
3 very-high performance clock (such as an audio sample clock), it could potentially  
4 achieve even higher resolution and lower jitter than the built-in clock/sequencer.

5 Module 446 operates as a hardware driver customized to the MIDI output  
6 device 450. Module 446 converts the information in the received data packets 350  
7 to a form specific to the output device 450. Different manufacturers can use  
8 different signaling techniques, so the exact manner in which module 446 operates  
9 will vary based on speakers 450 (and/or output module 448). Module 446 is  
10 coupled to an output module 448 which synthesizes the MIDI data into sound that  
11 can be played by speakers 450. Although illustrated in the software level, output  
12 module 448 may alternatively be implemented in the hardware level. By way of  
13 example, module 446 may be a MIDI output module which synthesizes MIDI  
14 messages into sound, a MIDI-to-waveform converter (often referred to as a  
15 software synthesizer), etc. In one implementation, output module 448 is included  
16 as part of a hardware driver corresponding to output device 450.

17 Module 446 is connected to allocator module 436. After the data for a data  
18 packet has been communicated to the output device 450, module 446 calls the  
19 PutMessage interface of the module it is connected to (allocator 436) to return the  
20 memory space used by the data packet to allocator 436 for re-allocation.

21 Another source of MIDI data illustrated in Fig. 9 is a user-mode  
22 application(s). A user-mode application can transmit MIDI data to unpacker  
23 module 452 in a buffer (such as buffer 380 of Fig. 7 or buffer 410 of Fig. 8).  
24 Analogous to packer module 442 discussed above, different unpacker modules can  
25 be used as unpacker module 452, (one being dedicated to legacy applications and

1 the other being dedicated to non-legacy applications), or alternatively a single  
2 dual-mode unpacker module may be used. Unpacker module 452 operates to  
3 convert the MIDI data in the received buffer into data packets 350, obtaining  
4 memory space from allocator module 436 for generation of the data packets 350.  
5 Unpacker module 452 is connected to sequencer module 444. Once a data packet  
6 350 is created, unpacker module 452 calls its programmed PutMessage interface to  
7 transmit the data packet 350 to sequencer module 444. Sequencer module 444,  
8 upon receipt of the data packet 350, operates as discussed above to either queue  
9 the data packet 350 or immediately transfer the data packet 350 to module 446.  
10 Because the unpacker 450 has done its job of converting the data stream from a  
11 large buffer into smaller individual data packets, these data packets can be easily  
12 sorted and interleaved with a data stream also entering the sequencer 444 – from  
13 the splitter 440 for example.

14 Fig. 10 is a block diagram illustrating another exemplary MIDI transform  
15 module graph 454 such as may be used in accordance with certain embodiments of  
16 the invention. Graph 454 of Fig. 10 is similar to graph 430 of Fig. 9, except that  
17 one or more additional modules 456 that perform various operations are added to  
18 graph 454 by graph builder 312 of Fig. 3. As illustrated, one or more of these  
19 additional modules 456 can be added in graph 454 in a variety of different  
20 locations, such as between modules 438 and 440, between modules 440 and 442,  
21 between modules 440 and 444, between modules 452 and 444, and/or between  
22 modules 444 and 446.

23 Fig. 11 is a flowchart illustrating an exemplary process for the operation of  
24 a module in a MIDI transform module graph in accordance with certain  
25 embodiments of the invention. In the illustrated example, the process of Fig. 11 is

1 implemented by a software module (e.g., module 326 of Fig. 3) executing on a  
2 computing device.

3 Initially, a data packet, including MIDI data (e.g., a data packet 350 of Fig.  
4 5) is received by the module (act 462). Upon receipt of the MIDI data, the module  
5 processes the MIDI data (act 464). The exact manner in which the data is  
6 processed is dependent on the particular module, as discussed above. Once  
7 processing is complete, the programmed PutMessage interface (which is on a  
8 different module) is called (act 468), forwarding the data packet to the next  
9 module in the graph.

10 Fig. 12 is a flowchart illustrating an exemplary process for the operation of  
11 a graph builder in accordance with certain embodiments of the invention. In the  
12 illustrated example, the process of Fig. 12 is carried out by a graph builder 312 of  
13 Fig. 3 implemented in software. Fig. 12 is discussed with additional reference to  
14 Fig. 3. Although a specific ordering of acts is illustrated in Fig. 12, the ordering of  
15 the acts can alternatively be re-arranged.

16 Initially, graph builder 312 receives a request to build a graph (act 472).  
17 This request may be for a new graph or alternatively to modify a currently existing  
18 graph. The user-mode application 310 that submits the request to build the graph  
19 includes an identification of the functionality that the graph should include. This  
20 functionality can include any of a wide variety of operations, including pitch bends,  
21 volume changes, aftertouch alterations, etc. The user-mode application also  
22 submits, if relevant, an ordering to the changes. By way of example, the  
23 application may indicate that the pitch bend should occur prior to or subsequent to  
24 some other alteration.



1 In response to the received request, graph builder 312 determines which  
2 graph modules are to be included based at least in part on the desired functionality  
3 identified in the request (act 474). Graph builder 312 is programmed with, or  
4 otherwise has access to, information identifying which modules correspond to  
5 which functionality. By way of example, a lookup table may be used that maps  
6 functionality to module identifiers. Graph builder 312 also automatically adds  
7 certain modules into the graph (if not already present). In one implementation, an  
8 allocator module is automatically inserted, an unpacker module is automatically  
9 inserted for each output path, and packer and capture sink modules are  
10 automatically inserted for each input path.

11 Graph builder 312 also determines the connections among the graph  
12 modules based at least in part on the desired functionality (and ordering, if any)  
13 included in the request (act 476). In one implementation, graph builder 312 is  
14 programmed with a set of rules regarding the building of graphs (e.g., which  
15 modules must or should, if possible, be prior to which other modules in the graph).  
16 Based on such a set of rules, the MIDI transform module graph can be constructed.

17 Graph builder 312 then initializes any needed graph modules (act 478).  
18 The manner in which graph modules are initialized can vary depending on the type  
19 of module. For example, pointers to the allocator module and reference clock may  
20 be passed to the module, other operating parameters may be passed to the module,  
21 etc.

22 Graph builder then adds any needed graph modules (as determined in act  
23 474) to the graph (act 480), and connects the graph modules using the connections  
24 determined in act 476 (act 482). If any modules need to be temporarily paused to  
25 perform the connections, graph builder 312 changes the state of such graph

1 modules to a stop state (act 484), which may involve transitioning between one or  
2 more intermediate states (e.g., pause and/or acquire states). The outputs for the  
3 added modules are connected first, and then the other modules are redirected to  
4 feed them, working in a direction "up" the graph from destination to source (act  
5 486). This reduces the chances that the graph would need to be stopped to insert  
6 modules. Once connected, any modules in the graph that are not already in a run  
7 state are started (e.g., set to a run state) (act 488), which may involve transitioning  
8 between one or more intermediate states (e.g., pause and/or acquire states).  
9 Alternatively, another component may set the modules in the graph to the run  
10 state, such as application 310. In one implementation, the component (e.g., graph  
11 builder 312) setting the nodes in the graph to the run state follows a particular  
12 ordering. By way of example, the component may begin setting modules to run  
13 state at a MIDI data source and follow that through to a destination, then repeat for  
14 additional paths in the graph (e.g., in graph 430 of Fig. 8, the starting of modules  
15 may be in the following order: modules 436, 434, 438, 440, 442, 444, 446, 452).  
16 Alternatively, certain modules may be in a "start first" category (e.g., allocator 436  
17 and sequencer 444 of Fig. 8).

18 In one implementation, graph builder 312 follows certain rules when  
19 adding or deleting items from the graph as well as when starting or stopping the  
20 graph. Reference is made herein to "merger" modules, branching modules, and  
21 branches within a graph. Merging is built-in to the interface described above, and  
22 a merger module refers to any module that has two or more other modules  
23 outputting to it (that is, two or more other modules calling its PutMessage  
24 interface). Graph builder 312 knows this information (who the mergers are),  
25 however the mergers themselves do not. A branching module refers to any module

1 from which two or more branches extend (that is, any module that duplicates (at  
2 least in part) data and forwards copies of the data to multiple modules). An  
3 example of a branching module is a splitter module. A branch refers to a string of  
4 modules leading to or from (but not including) a branching module or merger  
5 module, as well as a string of modules between (but not including) merger and  
6 branching modules.

7 When moving the graph from a lower state (e.g., stop) to a higher state  
8 (e.g., run), graph builder 312 first changes the state of the destination modules,  
9 then works its way toward the source modules. At places where the graph  
10 branches (e.g., splitter modules), all destination branches are changed before the  
11 branching module (e.g., splitter module) is changed. In this way, by the time the  
12 "spigot is turned on" at the source, the rest of the graph is in run state and ready to  
13 go.

14 When moving the graph from a higher state (e.g., run) to a lower state (e.g.,  
15 stop), the opposite tack is taken. First graph builder 312 stops the source(s), then  
16 continues stopping the modules as it progresses toward the destination module(s).  
17 In this way the "spigot is turned off" at the source(s) first, and the rest of the graph  
18 is given time for data to empty out and for the modules to "quiet" themselves. A  
19 module quieting itself refers to any residual data in the module being emptied out  
20 (e.g., an echo is passively allowed to die off, etc.). Quieting a module can also be  
21 actively accomplished by putting the running module into a lower state (e.g., the  
22 pause state) until it is no longer processing any residual data (which graph builder  
23 312 can determine, for example, by calling its GetParameters interface).

24 When a module is in stop state, the module fails any calls to the module's  
25 PutMessage interface. When the module is in the acquire state, the module

1 accepts PutMessage calls without failing them, but it does not forward messages  
2 onward. When the module is in the pause state, it accepts PutMessage calls and  
3 can work normally as long as it does not require the clock (if it needs a clock, then  
4 the pause state is treated the same as the acquire state). Clockless modules are  
5 considered "passive" modules that can operate fully during the "priming" sequence  
6 when the graph is in the pause state. Active modules only operate when in the run  
7 state. By way of example, splitter modules are passive, while sequencer modules,  
8 miniport streams, packer modules, and unpacker modules are active.

9 Different portions of a graph can be in different states. When a source is  
10 inactive, all modules on that same branch can be inactive as well. Generally, all  
11 the modules in a particular branch should be in the same state, including source  
12 and destination modules if they are on that branch. Typically, the splitter module  
13 is put in the same state as its input module. A merger module is put in the highest  
14 state (e.g., in the order stop, pause, acquire, run) of any of its input modules.

15 Graph builder 312 can insert modules to or delete modules from a graph  
16 "live" (while the graph is running). In one implementation, any module except  
17 miniport streams, packers, unpackers, capture sinks, and sequencers can be  
18 inserted to or deleted from the graph while the graph is running. If a module is to  
19 be added or deleted while the graph is running, care should be taken to ensure that  
20 no data is lost when making changes, and when deleting a module that the module  
21 is allowed to completely quiet itself before it is disconnected.

22 By way of example, when adding a module B between modules A and C,  
23 first the output of module B is connected to the input of module C (module C is  
24 still being fed by module A). Then, graph builder 312 switches the output of  
25 module A from module C to module B with a single ConnectOutput call. The

1 module synchronizes ConnectOutput calls with PutMessage calls, so  
2 accomplishing the graph change with a single ConnectOutput call ensures that no  
3 data packets are lost during the switchover. In the case of a branching module, all  
4 of its outputs are connected first, then its source is connected. When adding a  
5 module immediately previous to a merger module (where the additional module is  
6 intended to be common to both data paths), the additional module becomes the  
7 new merger module, and the item that was previously considered a merger module  
8 is no longer regarded as a merger module. In that case, the new merger module's  
9 output and the old merger module's input are connected first, then the old merger  
10 module's inputs are switched to the new merger module's inputs. If it is absolutely  
11 necessary that all of the merger module's inputs switch to the new merger at the  
12 same instant, then a special SetParams call should be made to each of the  
13 "upstream" input modules to set a timestamp for when the ConnectOutput should  
14 take place.

15 When deleting a module B from between modules A and C, first the output  
16 of module A is connected to the input of module C (module B is effectively  
17 bypassed at this time). Then, after module B empties and quiets itself (e.g., it  
18 might be an echo or other time-based effect), its output is reset to the allocator.  
19 Then module B can be safely destroyed (e.g., removed from the graph). When  
20 deleting a merger module, first its inputs are switched to the subsequent module  
21 (which becomes a merger module now), then after the old merger module quiets,  
22 its output is disconnected. When deleting a branching module, this is because an  
23 entire branch is no longer needed. In that case, the branching module output going  
24 to that branch is disconnected. If the branching module had more than two  
25 outputs, then the graph builder calls DisconnectOutput to disconnect that output

1 from the branching module's output list. At that point the subsequent modules in  
2 that branch can be safely destroyed. However, if the branching module had only  
3 two connected outputs, then the splitter module is no longer necessary. In that  
4 case, the splitter module is bypassed (the previous module's output is connected to  
5 the subsequent module's input), then after the splitter module quiets it is  
6 disconnected and destroyed.

### 7 8 **Transform Modules**

9 Specific examples of modules that can be included in a MIDI transform  
10 module graph (such as graph 430 of Fig. 9, graph 454 of Fig. 10, or graph 314 of  
11 Fig. 3) are described above. Various additional modules can also be included in a  
12 MIDI transform module graph, allowing user-mode applications to generate a  
13 wide variety of audio effects. Furthermore, as graph builder 312 of Fig. 3 allows  
14 the MIDI transform module graph to be readily changed, the functionality of the  
15 MIDI transform module graph can be changed to include new modules as they are  
16 developed.

17 Fig. 13 is a block diagram illustrating an exemplary set of additional  
18 transform modules that can be made added to a module graph in accordance with  
19 certain embodiments of the invention. In one implementation, the set of transform  
20 modules 520 is included in module library 330. These exemplary additional  
21 modules 520 are described in more detail below.

22 These additional modules include the four common interfaces discussed  
23 above (SetState, PutMessage, ConnectOutput, and DisconnectOutput). For  
24 modules that use parameters (e.g., specific channel numbers, specific offsets, etc.),  
25 these parameters can be set via a SetParameters interface, or alternatively multiple

1 versions of the modules can be generated with pre-programmed parameters (which  
2 of the modules to include in the graph is then dependent on which parameters  
3 should be used).

4 In the illustrated example, graph builder 312 of Fig. 3 passes any necessary  
5 parameters to the modules during initialization. Which parameters are to be  
6 passed to a module are received by graph builder 312 from application 310. By  
7 way of example, application 310 may indicate that a particular channel is to be  
8 muted (e.g., due to its programming, due to inputs from a user via a user interface,  
9 etc.).

10 The additional modules described below may also include a GetParameters  
11 interface, via which graph builder 312 (or alternatively application 310 or another  
12 module 326) may obtain information from the modules. This information will  
13 vary, depending on the module. By way of example, the parameters used by a  
14 module (whether set via a SetParameters interface or pre-programmed) can be  
15 obtained by the GetParameters interface, or information being gathered (e.g.,  
16 about the graph) or maintained by a module may be obtained by the  
17 GetParameters interface.

18 In one implementation, each of these additional modules is passed a pointer  
19 to an allocator module as well as a reference clock, as discussed above.  
20 Alternatively, one or more of the additional modules may not be passed the pointer  
21 to the allocator module and/or the reference clock.

22 For ease of explanation, the additional transform modules are discussed  
23 herein with reference to operating on data included within a data packet (e.g., data  
24 packet 350 of Fig. 6). It is to be appreciated that these transform modules may  
25 also operate on data that is contained within a chain of data packets pointed to by a

1 particular data packet 350, or on audio data (e.g., messages 345 of Fig. 5) included  
2 in a data buffer pointed to by a particular data packet 350.

3 It is to be appreciated that, when handling packet chains, if one or more  
4 events are removed from the chain by a module then the next event portion 366 of  
5 a preceding event (and possibly the event chain pointer 372 of data packet 350)  
6 may need to be updated to accurately identify the next event in the chain. For  
7 example, if an event chain includes three events and the second event is removed  
8 from the chain, then the next event portion 366 of the first event is modified to  
9 identify the last event in the chain (rather than the second event which it  
10 previously identified).

11 The sequencer, splitter, capture sink, and allocator modules are discussed  
12 above in greater detail. A sequencer module does not change the data itself, but it  
13 does reorder the data by timestamp and delay forwarding the message on to the  
14 next module in the graph until the appropriate time. A splitter module creates one  
15 or more additional data packets virtually identical to the input data packets  
16 (obtaining additional data packets from an allocator module to do so). A capture  
17 sink module takes audio data that is either parsed or unparsed, and emits a parsed  
18 audio data stream. An allocator module obtains memory from a memory manager  
19 and makes portions of the obtained memory available for audio data.

20 Unpacker. Unpacker modules, in addition to those discussed above, can  
21 also be included in a MIDI transform module graph. Unpacker modules operate to  
22 receive data into the graph from a user-mode application, converting the MIDI  
23 data received in the user-mode application format into data packets 350 (Fig. 6)  
24 for communicating to other modules in the graph. Additional unpacker modules,  
25



1 supporting any of a wide variety of user-mode application specific formats, can be  
2 included in the graph.

3 Packer. Packer modules, in addition to those discussed above, can also be  
4 included in a MIDI transform module graph. Packer modules operate to output  
5 MIDI data from the graph to a user-mode application, converting the MIDI data  
6 from the data packets 350 into a user-mode application specific format.  
7 Additional packer modules, supporting any of a wide variety of user-mode  
8 application specific formats, can be included in the graph.

9 Feeder In. A Feeder In module operates to convert MIDI data received in  
10 from a software component that is not aware of the data formats and protocols  
11 used in a module graph (e.g., graph 314 of Fig. 3) into data packets 350. Such  
12 components are typically referred to as "legacy" components, and include, for  
13 example, older hardware miniport drivers. Different Feeder In modules can be  
14 used that are specific to the particular hardware drivers they are receiving the  
15 MIDI data from. The exact manner in which the Feeder In modules operate will  
16 vary, depending on what actions are necessary to convert the received MIDI data  
17 to the data packets 350.

18 Feeder Out. A Feeder Out module operates to convert MIDI data in data  
19 packets 350 into the format expected by a particular legacy component (e.g., older  
20 hardware miniport driver) that is not aware of the data formats and protocols used  
21 in a module graph (e.g., graph 314 of Fig. 3). Different Feeder Out modules can  
22 be used that are specific to the particular hardware drivers they are sending the  
23 MIDI data to. The exact manner in which the Feeder Out modules operate will  
24 vary, depending on what actions are necessary to convert the MIDI data in the data  
25 packets 350 into the format expected by the corresponding hardware driver.

1        Channel Mute. A Channel Mute module operates to mute one or more  
2 MIDI channel(s) it has set as a parameter. A Channel Mute module can be  
3 channel-only or channel and group combined. As discussed above, the MIDI  
4 standard allows for multiple different channels (encoded in status byte 346 of  
5 message 345 of Fig. 5). The data packet 350, however, allows for multiple  
6 channel groups (identified in channel group portion 358). The parameter(s) for a  
7 Channel Mute module can identify a particular channel (e.g., channel number five,  
8 regardless of which channel group it is in) or a combination of channel and group  
9 number (e.g., channel number five in channel group number 692).

10        Upon receipt of a data packet 350, the channel mute module checks which  
11 channel the data packet 350 corresponds to. The channel mute module compares  
12 its parameter(s) to the channel that data packet 350 corresponds to. If the channel  
13 matches at least one of the parameters (e.g., is the same as at least one of the  
14 parameters), then data packet 350 is forwarded to the allocator module for re-  
15 allocation of the memory space. The data is not forwarded for further audio  
16 processing, effectively muting the channel. However, if the channel does not  
17 match at least one of the parameters, then data packet 350 is forwarded on for  
18 further audio processing.

19        Channel Solo. A Channel Solo module operates to pass through only a  
20 selected channel(s). A Channel Solo module operates similarly to a Channel Mute  
21 module, comparing the parameter(s) to a channel that data packet 350 corresponds  
22 to. However, only those packets 350 that correspond to a channel(s) that matches  
23 at least one of the parameter(s) are forwarded for further audio processing; packets  
24 350 that correspond to a channel that does not match at least one of the parameters  
25 are forwarded to the allocator module for re-allocation of the memory space.

1        Channel Route. A Channel Route module operates to alter a particular  
2 channel. A Channel Route module typically includes one source channel and one  
3 destination channel as a parameter. The channel that a data packet 350  
4 corresponds to is compared to the source channel parameter, analogous to a  
5 Channel Mute module discussed above. However, if a match is found, then the  
6 channel number is changed to the destination channel parameter (that is, status  
7 byte 346 is altered to encode the destination channel number rather than the source  
8 channel number). Data packets 350 received by a Channel Route module are  
9 forwarded on to the next module in the graph for further audio processing  
10 (whatever module(s) the Channel Route module is connected to) regardless of  
11 whether the channel number has been changed.

12        Channel Route/Map. A Channel Route/Map module operates to alter  
13 multiple channels. A Channel Route/Map module is similar to a Channel Route  
14 module, except that a Channel Route/Map module maps multiple source channels  
15 to one or more different destination channels. In one implementation, this is a 1 to  
16 1 mapping (each source channel is routed to a different destination channel). The  
17 source and destination channel mappings are a parameter of the Channel  
18 Route/Map module. In one implementation, a Channel Route/Map module can re-  
19 route up to sixteen different source channels (e.g., the number of channels  
20 supported by the MIDI standard). Data packets 350 received by a Channel  
21 Route/Map module are forwarded on to the next module in the graph for further  
22 audio processing (whatever module(s) the Channel Route/Map module is  
23 connected to) regardless of whether the channel number has been changed.

24        Channel Map. A Channel Map module operates to provide a general case  
25 of channel mapping and routing, allowing any one or more of the sixteen possible

1 channels to be routed to any one or more of the sixteen possible channels. This  
2 mapping can be one to one, one to many, or many to one. Data packets 350  
3 received by a Channel Map module (as well as any data packets generated by a  
4 Channel Map module) are forwarded on to the next module in the graph for further  
5 audio processing (whatever module(s) the Channel Map module is connected to)  
6 regardless of whether the channel number has been changed.

7 In one implementation, a Channel Map module includes a 16x16 matrix as  
8 a parameter. Fig. 14 illustrates an exemplary matrix 540 for use in a Channel Map  
9 module in accordance with certain embodiments of the invention. Channel inputs  
10 (source channels) are identified along the Y-axis and channel outputs (destination  
11 channels) are identified along the X-axis. A value of one in the matrix indicates  
12 that the corresponding source channel is to be changed to the corresponding  
13 destination channel, while a value of zero in the matrix indicates that the  
14 corresponding source channel is not to be changed.

15 In the illustrated matrix 540, if the source channel is 2, 4, 5, 7, 8, 9, 10, 12,  
16 13, 14, 15, or 16, then no change is made to the channel. If the source channel is  
17 1, then the destination channel is 5, so the channel number is changed to 5. If the  
18 source channel is 3, then the destination channels are 1, 8, and 15. The Channel  
19 Map module can either keep the data packet with the source channel of 3 and  
20 generate new packets with channels of 1, 8, and 15, or alternatively change the  
21 data packet with the source channel of 3 to one of the channels 1, 8, or 15 and then  
22 create new packets for the remaining two destination channels. If any new packets  
23 are to be created, the Channel Map module obtains new data packets from the  
24 allocator module (via its GetMessage interface). If the source channel is 6, then  
25 the channel number is changed to 5, and if the source channel is 11, then the

1 channel number is changed to 14. It should be noted that any packets having a  
2 corresponding channel number of either 1 or 6 will have the channel number  
3 changed to 5 by the Channel Map module, resulting in a "many to one" mapping.

4 Channel Group Mute. A Channel Group Mute module operates to mute  
5 channel groups. A Channel Group Mute module operates similar to a Channel  
6 Mute module, except that a Channel Group Mute module operates to mute groups  
7 of channels rather than individual channels. One or more channel groups can be  
8 set as the mute parameter(s). The channel group identified in channel group  
9 portion 358 of a packet 350 is compared to the parameter(s). If the channel group  
10 from the packet matches at least one of the parameter(s), then packet 350 is  
11 forwarded to the allocator module for re-allocation of the memory space;  
12 otherwise, the packet 350 is forwarded on for further audio processing.

13 Channel Group Solo. A Channel Group Solo module operates to delete all  
14 except selected channel groups. A Channel Group Solo module operates similarly  
15 to a Channel Group Mute module, comparing the parameter(s) to a channel group  
16 that data packet 350 corresponds to. However, only those packets 350 that  
17 correspond to a channel group(s) that matches at least one of the parameter(s) are  
18 forwarded for further audio processing; packets 350 that correspond to a channel  
19 group that does not match the parameter are forwarded to the allocator module for  
20 re-allocation of the memory space.

21 Channel Group Route. A Channel Group Route module operates to route  
22 groups of channels. A Channel Group Route module operates similar to a Channel  
23 Route module, except that a Channel Group Route module operates to alter a  
24 particular group of channels rather than individual channels. One or more channel  
25 groups can be set as the route parameter(s). A Channel Group Route module

1 typically includes one source channel group and one destination channel group as  
2 parameters. The channel group that a data packet 350 corresponds to is compared  
3 to the source channel group parameter, analogous to the Channel Route module  
4 discussed above. However, if a match is found, then the channel group number is  
5 changed to the destination channel group parameter (that is, channel group portion  
6 358 is altered to include the destination channel group number rather than the  
7 source channel group number). Data packets 350 received by a channel group  
8 route module are forwarded on for further audio processing regardless of whether  
9 the channel group number has been changed.

10 Channel Group Map. A Channel Group Map module operates to alter  
11 multiple channel groups. A Channel Group Map module is similar to a Channel  
12 Group Route module, except that a Channel Group Map module maps multiple  
13 source channel groups to one or more different destination channel groups. In one  
14 implementation, this is a 1 to 1 mapping (each source channel group is routed to a  
15 different destination channel group). The source and destination channel group  
16 mappings, as well as the number of such mappings, are parameters of a Channel  
17 Group Map module.

18 Message Filter. A Message Filter module operates to allow certain types of  
19 messages through while other types of messages are blocked. According to the  
20 MIDI standard, there are 128 different status byte possibilities (allowing for 128  
21 different types of messages). In one implementation, a 128-bit buffer is used as a  
22 "bit mask" to allow selected ones of these 128 different types of messages through  
23 while others are blocked. This 128-bit bit mask buffer is the parameter for a  
24 Message Filter module. Each of the 128 different message types is assigned a  
25 number (this is inherent in the use of 7 bits to indicate message type, as  $2^7 = 128$ ).

1 This number is then compared to the corresponding bit in the bit mask buffer. By  
2 way of example, if the 7 bits of the status byte that indicate the message type are  
3 0100100 (which equals decimal 36), then the message filter module would check  
4 whether the 36<sup>th</sup> bit of the bit mask buffer is set (e.g., a value of one). If the 36<sup>th</sup>  
5 bit is set, then the message is allowed to pass through (that is, it is forwarded on  
6 for further audio processing). However, if the 36<sup>th</sup> bit is not set (e.g., a value of  
7 zero), then the message is blocked (that is, it is forwarded to the allocator module  
8 so that the memory space can be re-allocated).

9 Note Offset. A Note Offset module operates to transpose note by a given  
10 offset value. A signed offset value (e.g., a 7-bit value) is a parameter for a Note  
11 Offset module, as well as the channel(s) (and/or channel group(s)) that are to have  
12 their notes transposed. When a data packet 350 is received, a check is made as to  
13 whether the channel(s) and or channel group(s) corresponding to the message  
14 included in data portion 368 of packet 350 match at least one of the parameters. If  
15 there is a match, then the Note Offset module alters the value of the note by the  
16 offset value. This alteration can be performed either with or without rollover. For  
17 example, assuming there are 128 notes, that the note value for the message is 126,  
18 and that the offset is +4, the alteration could be without rollover (e.g., change the  
19 note value to 128), or with rollover (e.g., change the note value to 2).

20 Data packets 350 received by a Note Offset module are forwarded on to the  
21 next module in the graph for further audio processing regardless of whether the  
22 note value has been changed.

23 Note Map Curve. A Note Map Curve module operates to allow individual  
24 transposition of notes. An input note to output note mapping table is used as a  
25 parameter for a Note Map Curve module, the table identifying what each of the

1 input notes is to be mapped to. When a data packet 350 is received, the note  
2 identified in data portion 368 is compared to the mapping table. The mapping  
3 table identifies an output note value, and the Note Map Curve module changes the  
4 value of the note identified in data portion 368 to the output note value.

5 The MIDI standard supports 128 different note values. In one  
6 implementation, the mapping table is a table including 128 entries that are each 7  
7 bits. Each of the 128 entries corresponds to one of the 128 different notes (e.g.,  
8 using the 7 bits that are used to represent the note value), and the corresponding  
9 entry includes a 7-bit value of what the note value should be mapped to.

10 Data packets 350 received by a Note Map Curve module are forwarded on  
11 to the next module in the graph for further audio processing regardless of whether  
12 the note value has been changed.

13 Note Palette Solo/Mute. A Note Palette Solo/Mute module operates to  
14 allow certain notes through for further audio processing while other notes are  
15 blocked. According to the MIDI standard, there are 128 different notes. In one  
16 implementation, a 128-bit buffer is used as a bit mask to allow selected ones of  
17 these 128 different notes through while others are blocked. This 128-bit bit mask  
18 buffer is the parameter for a Note Palette Solo/Mute module. Each of the 128  
19 different notes is assigned a number (this is inherent in the use of 7 bits to indicate  
20 message type, as  $2^7 = 128$ ). This number is then compared to the corresponding  
21 bit in the bit mask buffer. By way of example, if the 7 bits indicating the value of  
22 the note are 1101011 (which equals decimal 107), then a Note Palette Solo/Mute  
23 module checks whether the 107<sup>th</sup> bit of the bit mask buffer were set (e.g., a value  
24 of one). If the 107<sup>th</sup> bit is set, then the Note Palette Solo/Mute module allows the  
25 packet corresponding to the note to pass through (that is, the packet including the



1 note message is forwarded on for further audio processing in the graph).  
2 However, if the 107<sup>th</sup> bit is not set (e.g., a value of zero), then the Note Palette  
3 Solo/Mute module blocks the note (that is, the packet including the note message  
4 is forwarded to the allocator module so that the memory space can be re-  
5 allocated).

6 Note Palette Adjuster. A Note Palette Adjuster module operates to snap  
7 "incorrect" notes to the closest valid note. A Note Palette Adjuster module  
8 includes, as a parameter, a bit mask analogous to that of a Note Palette Solo/Mute  
9 module. If the bit in the bit mask corresponding to a note is set, then the Note  
10 Palette Adjuster module allows the packet corresponding to the note to pass  
11 through (that is, the packet including the note message is forwarded on for further  
12 audio processing in the graph). However, if the bit in the bit mask corresponding  
13 to the note is not set, then the note is "incorrect" and the Note Palette Adjuster  
14 module changes the note value to be the closest "valid" value (that is, the closest  
15 note value for which the corresponding bit in the bit mask is set). If two notes are  
16 the same distance to the incorrect note, then the Note Palette Adjuster module uses  
17 a "tie-breaking" process to select the closest note (e.g., always go to the higher  
18 note, always go to the lower note, go the same direction (higher or lower) as was  
19 used for the previous incorrect note, etc.).

20 Data packets 350 received by a Note Palette Adjuster module are forwarded  
21 on to the next module in the graph for further audio processing regardless of  
22 whether the note value has been changed.

23 Velocity Offset. A Velocity Offset module operates to alter the velocity of  
24 notes by a given offset value. A signed offset value (e.g., a 7-bit value) is a  
25 parameter for a Velocity Offset module. Additional parameters optionally include

1 the note(s), channel(s), and/or channel group(s) that will have their velocities  
2 altered. When a data packet 350 is received, the Velocity Offset module compares  
3 the note(s), channel(s), and channel group(s) (if any) parameters to the note(s),  
4 channel(s), and channel group(s) corresponding to the message included in data  
5 portion 368 of packet 350 to determine whether there is a match (e.g., if they are  
6 the same). If there is a match (or if there are no such parameters), then the  
7 Velocity Offset module alters the velocity value for the message included in data  
8 portion 368 of packet 350 (e.g., as encoded in status byte 346 of message 345 of  
9 Fig. 5) by the offset value. This alteration can be performed either with or without  
10 rollover.

11 Data packets 350 received by a Velocity Offset module are forwarded on to  
12 the next module in the graph for further audio processing regardless of whether the  
13 velocity value has been changed.

14 Velocity Map Curve. A Velocity Map Curve module operates to allow  
15 individual velocity alterations. An input velocity to output velocity mapping table  
16 is used as a parameter for the Velocity Map Curve module, the table identifying  
17 what each of the input velocities is to be mapped to. When a data packet 350 is  
18 received, the velocity identified in data portion 368 (e.g., as encoded in status byte  
19 346 of message 345 of Fig. 5) is compared to the mapping table. The mapping  
20 table identifies an output velocity value, and the Velocity Map Curve module  
21 changes the value of the velocity identified in data portion 368 to the output  
22 velocity value from the table.

23 The MIDI standard supports 128 different velocity values. In one  
24 implementation, the mapping table is a table including 128 entries that are each 7  
25 bits (analogous to that of the Note Map Curve module discussed above). Each of

1 the 128 entries corresponds to one of the 128 different velocity values (e.g., using  
2 the 7 bits that are used to represent the velocity value), and the corresponding  
3 entry includes a 7-bit value of what the velocity value should be mapped to.

4 Data packets 350 received by a Velocity Map Curve module are forwarded  
5 on to the next module in the graph for further audio processing regardless of  
6 whether the velocity value has been changed.

7 Note and Velocity Map Curve. A Note and Velocity Map Curve module  
8 operates to allow combined note and velocity alterations based on both the input  
9 note and velocity values. A parameter for the Note and Velocity Map Curve  
10 module is a mapping of input note and velocity to output note and velocity. In one  
11 implementation, this mapping is a table including 16,384 entries (one entry for  
12 each possible note and velocity combination, assuming 128 possible note values  
13 and 128 possible velocity values) that are each 14-bits (7 bits indicating the new  
14 note value and 7 bits indicating the new velocity value). When a data packet 350  
15 is received, the velocity and note identified in data portion 368 (e.g., as encoded in  
16 status byte 346 of message 345 of Fig. 5) is compared to the mapping table. The  
17 mapping table identifies an output velocity value and an output note value, and the  
18 Note and Velocity Map Curve module changes the value of the velocity identified  
19 in data portion 368 to the output velocity value from the table.

20 The Note and Velocity Map Curve module may generate a new data packet  
21 rather than change the value of the note (this can be determined, for example, the  
22 setting of an additional bit in each entry of the mapping table). The input data  
23 packet would remain unchanged, and a new data packet would be generated that is  
24 a duplicate of the input data packet except that the new data packet includes the  
25 note and velocity values from the mapping table.

1 Data packets 350 received by a Note and Velocity Map Curve module are  
2 forwarded on to the next module in the graph for further audio processing  
3 regardless of whether the note and/or velocity values have been changed.

4 Time Offset. A Time Offset module operates to alter the presentation time  
5 of notes by a given offset value. A signed offset value (e.g., an 8-byte value) is a  
6 parameter for a Time Offset module. In one implementation, the offset value is in  
7 the same units as are used for presentation time portion 362 of data packet 350  
8 (e.g., 100 ns units). Additional parameters optionally include the note(s),  
9 channel(s), and/or channel group(s) that will have their presentation times altered.  
10 When a data packet 350 is received, the Time Offset module compares the note(s),  
11 channel(s), and channel group(s) (if any) parameters to the note(s), channel(s), and  
12 channel group(s) corresponding to the message included in data portion 368 of  
13 packet 350 to determine whether there is a match (e.g., if they are the same). If  
14 there is a match (or if there are no such parameters), then the Time Offset module  
15 alters the presentation time in portion 362 of packet 350 by the offset value. This  
16 alteration can be performed either with or without rollover.

17 Data packets 350 received by a Time Offset module are forwarded on to the  
18 next module in the graph for further audio processing regardless of whether the  
19 presentation time value has been changed.

20 Time Palette. A Time Palette module operates to alter the presentation  
21 times of notes. A grid (e.g., mapping input presentation times to output  
22 presentation times) or multiplier is used as a parameter to a Time Palette module,  
23 and optionally an offset as well. Additional parameters optionally include the  
24 note(s), channel(s), and/or channel group(s) that will have their presentation times  
25 altered. When a data packet 350 is received, the Time Palette module compares

1 the note(s), channel(s), and channel group(s) (if any) parameters to the note(s),  
2 channel(s), and channel group(s) corresponding to the message included in data  
3 portion 368 of packet 350 to determine whether there is a match (e.g., if they are  
4 the same). If there is a match (or if there are no such parameters), then the Time  
5 Palette module alters the presentation time in portion 362 of packet 350 to be that  
6 of the closest multiplier (or grid entry) – that is, the presentation time is "snapped"  
7 to the closest multiplier (or grid entry). The optional offset parameter is used by  
8 the Time Palette module to indicate how the multiplier is to be applied. For  
9 example, if the multiplier is ten and the offset is two, then the presentation times  
10 are changed to the closest of 2, 12, 22, 32, 42, 52, 62, etc. This "snapping" process  
11 is referred to as a quantization process.

12 Alternatively, rather than snapping to the closest multiplier (or grid entry),  
13 the presentation times could be snapped closer to the closest multiplier (or grid  
14 entry). How close the presentation times are snapped can be an additional  
15 parameter for the Time Palette module (e.g., 2 ns closer, 50% closer, etc.).

16 The Time Palette module can also perform an anti-quantization process. In  
17 an anti-quantization process, the Time Palette module uses an additional parameter  
18 that indicates the maximum value that presentation times of notes should be  
19 moved. The Time Palette module then uses an algorithm to determine, based on  
20 the maximum value parameter, how much the presentation time should be moved.  
21 This algorithm could be, for example, a random number generator, or alternatively  
22 an algorithm to identify the closest multiplier (or grid entry) to be snapped to and  
23 then adding (or subtracting) a particular amount (e.g., a random value) to that  
24 "snap" point.  
25

1 Time palette modules can also operate to alter the rhythmic feel of music,  
2 such as to include a "swing" feel to the music. Two additional parameters are  
3 included for the Time Palette module to introduce swing: a subdivision value and  
4 a desired balance. The subdivision value indicates the amount of time (e.g., in 100  
5 ns units) between beats. The desired balance indicates how notes within this  
6 subdivision should be altered. This in effect is creating a virtual midpoint between  
7 beats that is not necessarily exactly 50% between the beats, and the balance  
8 parameter determines exactly how close to either side that subbeat occurs. The  
9 Time Palette module does not change any note that occurs on the beat (e.g., a  
10 multiplier of the subdivision amount). However, the Time Palette module alters  
11 any note(s) that occurs between the beat by "pushing" them out by an amount  
12 based on the desired balance, either toward the beat or toward the new "virtual  
13 half-beat". For example, if the subdivision amount is 100 then the subbeat value  
14 would be 50 (a beat is still 100). However, if the desired balance were 65, then the  
15 presentation times of notes between the beat are incremented so that half of the  
16 notes are between 0 and 65, and the other half are between 65 and 100. Notes that  
17 came in with timestamps of 0, 50, 100, 150, etc. would be changed to 0, 65, 100,  
18 165, etc.

19 Pitch Bend. A Pitch Bend module operates to bend the pitch for messages  
20 by a given offset value. A signed offset value (e.g., a 7-bit value) is a parameter  
21 for a Pitch Bend module. Additional parameters optionally include the note(s),  
22 channel(s), and/or channel group(s) that will have their pitches altered. When a  
23 data packet 350 is received (in one implementation, only when a data packet 350  
24 including a "pitch bend" type message is received), the Pitch Bend module  
25 compares the note(s), channel(s), and channel group(s) (if any) parameters to the

1 note(s), channel(s), and channel group(s) corresponding to the message included in  
2 data portion 368 of packet 350 to determine whether there is a match (e.g., if they  
3 are the same). If there is a match (or if there are no such parameters), then the  
4 Pitch Bend module alters the pitch value included in the message included in data  
5 portion 368 of packet 350 (e.g., encoded in data portion 347 of message 345 of  
6 Fig. 5) by the offset value. This alteration can be performed either with or without  
7 rollover.

8 Data packets 350 received by a Pitch Bend module are forwarded on to the  
9 next module in the graph for further audio processing regardless of whether the  
10 pitch value has been changed.

11 Variable Detune. A Variable Detune module operates to alter the pitch of  
12 (detune) music by a variable offset value. Parameters for a Variable Detune  
13 include a signed offset value (e.g., a 7-bit value) and a frequency indicating how  
14 fast over time the pitch is to be altered (e.g., the pitch should be altered from zero  
15 to 50 over a period of three seconds). Additional parameters optionally include  
16 the note(s), channel(s), and/or channel group(s) that will have their pitch values  
17 altered. When a data packet 350 is received (in one implementation, only when a  
18 data packet 350 including a "pitch bend" type message is received), the Variable  
19 Detune compares the note(s), channel(s), and channel group(s) (if any) parameters  
20 to the note(s), channel(s), and channel group(s) corresponding to the message  
21 included in data portion 368 of packet 350 to determine whether there is a match  
22 (e.g., if they are the same). If there is a match (or if there are no such parameters),  
23 then the Variable Detune alters the pitch value for the message included in data  
24 portion 368 of packet 350 (e.g., encoded in data portion 347 of message 345 of  
25 Fig. 5) by an amount based on the presentation time indicated in portion 362 of

1 packet 350 (or alternatively the current reference clock time) and the parameters.

2 This alteration can be performed either with or without rollover.

3         Given the offset and frequency parameters, the amount to alter the pitch  
4 value can be readily determined. Following the example above, the three second  
5 period of time can be broken into 50 equal portions, each assigned a value of one  
6 through 50 in temporal order. The assigned value to each portion is used to alter  
7 the pitch of any note with a presentation time corresponding to that portion. In  
8 one implementation, the offset and frequency parameters define an approximately  
9 sinusoidal waveform. In the above example, the waveform would start at zero, go  
10 to 50 over the first three seconds, then drop to zero over the next three seconds,  
11 then drop to negative 50 over the next three seconds, and then return from  
12 negative 50 to zero over the next three seconds, and then repeat (resulting in a  
13 period of 12 seconds).

14         Data packets 350 received by a Variable Detune module are forwarded on  
15 to the next module in the graph for further audio processing regardless of whether  
16 the pitch value has been changed.

17         Echo. An Echo module operates to generate an echo for notes. Time and  
18 velocity offsets are both parameters for the Echo module. Additional parameters  
19 optionally include the note(s), channel(s), and/or channel group(s) to be echoed.  
20 When a data packet 350 is received, the Echo module compares the note(s),  
21 channel(s), and channel group(s) (if any) parameters to the note(s), channel(s), and  
22 channel group(s) corresponding to the message included in data portion 368 of  
23 packet 350 to determine whether there is a match (e.g., if they are the same). If  
24 there is a match (or if there are no such parameters), then the Echo module obtains  
25 an additional data packet from the allocator module and copies the content of data



1 packet 350 into it, except that the velocity and presentation time of the new packet  
2 are altered based on the parameters. The time offset parameter indicates how  
3 much time is to be added to the presentation time of the new packet, and the  
4 velocity offset parameter indicates how much the velocity value of the message  
5 included in data portion 368 (e.g., encoded in status byte 346 of message 346 of  
6 Fig. 5) is to be reduced.

7 The echo module may also create multiple additional packets for a single  
8 packet that is being echoed, providing a series of packets with messages having  
9 continually reduced velocities and later presentation times. Each data packet in  
10 this series would differ from the previous packet in velocity and presentation time  
11 by an amount equal to the velocity and time offsets, respectively. Additional  
12 packets could be created until the velocity value drops below a threshold level  
13 (e.g., a fixed number or a percentage of the original velocity value), or a threshold  
14 number of additional packets have been created.

15 In one implementation, the Echo module forwards on the main message and  
16 feeds a copy of the data packet (after “weakening” it) to itself (e.g., either  
17 internally or via its PutMessage interface). This continues recursively until the  
18 incoming message is too weak to warrant an additional loop (back to the Echo  
19 module). In another implementation, all the resultant messages are computed at  
20 once and sent out immediately.

21 Additionally, a note delta may also be included as a parameter for an Echo  
22 module. The Echo module uses the note delta parameter to alter the note value of  
23 the message corresponding to the packet (in addition to altering the velocity and  
24 presentation time values). This results in an echo that changes in note as well as  
25 velocity (e.g., with notes spiraling upward or downward).

1 Alternatively, variable changes could be made to any of the velocity offset,  
2 note offset, or time offset values, resulting in a more random echo.

3 Data packets 350 received by an Echo module are forwarded on to the next  
4 module in the graph for further audio processing regardless of whether any Echo  
5 packets have been created.

6 Profile System Performance. A Profile System Performance module  
7 operates to monitor the system performance (e.g., with respect to jitter). Upon  
8 receipt of a data packet 350, a Profile System Performance module checks the  
9 presentation time 362 of the packet 350 and compares it to the current reference  
10 clock time. The Profile System Performance module records the difference and  
11 forwards the packet 350 to the next module in the graph. The Profile System  
12 Performance module maintains the recorded deltas and passes them to a requesting  
13 component (e.g., graph builder 312), such as in response to a call by graph builder  
14 312 to the GetParameters interface of the Profile System Performance module.

15 It is to be appreciated that the accuracy of the profile system performance  
16 module can be improved by locating it within the graph close to the rendering of  
17 the data (e.g., just prior to the passing of data packets 350 to module 446 of Fig.  
18 8).

19 Data packets 350 received by a Profile System Performance module are  
20 forwarded on to the next module in the graph for further audio processing  
21 regardless of whether any values have been recorded by the Profile System  
22 Performance module.

1 **Conclusion**

2       Although the description above uses language that is specific to structural  
3 features and/or methodological acts, it is to be understood that the invention  
4 defined in the appended claims is not limited to the specific features or acts  
5 described. Rather, the specific features and acts are disclosed as exemplary forms  
6 of implementing the invention.  
7  
8  
9  
10  
11  
12  
13  
14  
15  
16  
17  
18  
19  
20  
21  
22  
23  
24  
25